

Express Mail #ET029404372US

**APPLICATION
FOR
UNITED STATES LETTERS PATENT**

APPLICANT NAME: Michelle D. Campbell *et al.*

TITLE: GENERAL BUSINESS REPORT GENERATION

DOCKET NO.: END920010033US1

INTERNATIONAL BUSINESS MACHINES CORPORATION

GENERAL BUSINESS REPORT GENERATION

Background of the Invention

1. Technical Field

The present invention relates to a system and associated method for generating a report by
5 a reporting tool of the online financial software known as Systems Applications and Products
(SAP).

2. Related Art

The online financial software known as Systems Applications and Products (SAP)
includes the business data warehouse modules of Logistic Information System (LIS) and Open
10 Information Warehouse (OIW), and an associated reporting module called Executive Information
System (EIS) that receives and operates upon data from LIS and OIW. The SAP software is
owned by the SAP company in Germany. Unfortunately, data processing using EIS in
conjunction with the LIS and OIW modules for generating reports based on large volumes of data
is inefficient and prohibitively time consuming. Accordingly, there is a need for a time-efficient
15 method and system for generating reports within SAP.

Summary of the Invention

The present invention provides a system for generating a report by a reporting tool of a
SAP business information system using data included within an Aspect file, said system

comprising a non-SAP bridge program adapted to generate the Aspect file through use of data derived from a dataset and to transmit the Aspect file to the SAP business information system.

The present invention provides a system for generating a report by a reporting tool of a SAP business information system using data included within an Aspect file having rollup records, said system comprising a non-SAP bridge program adapted to generate the Aspect file through use of data derived from a dataset and to transmit the Aspect file to the SAP business information system, said dataset having a keygroup, wherein to generate the Aspect file includes to roll up a portion of the dataset with respect to the keygroup, wherein each rollup record has a rollup field and a quantity field, wherein the rollup field stores a rollup keyvalue of the keygroup, and wherein the quantity field stores the number of dataset records that have the same rollup keyvalue.

The present invention provides a method for generating a report by a reporting tool of a SAP business information system using data included within an Aspect file, said method comprising executing a non-SAP bridge program, said executing including: generating the Aspect file through use of data derived from a dataset; and transmitting the Aspect file to the SAP business information system.

The present invention provides a method for generating a report by a reporting tool of a SAP business information system using data included within an Aspect file having rollup records, said method comprising:

providing a dataset having a keygroup; and executing a non-SAP bridge program, including generating the Aspect file, said

generating comprising rolling up a portion of the dataset with respect to the keygroup, wherein each rollup record has a rollup field and a quantity field, wherein the rollup field stores a rollup keyvalue of the keygroup, and wherein the quantity field stores the number of dataset records that have the same rollup keyvalue.

5 The present invention provides a time-efficient method and system for generating reports

by the SAP Executive Information System (EIS). The present invention is relevant to any application that uses EIS, including applications involving procurement data, such as purchase order data and invoice data. Other applications may involve, *inter alia*, financial situations, human resources (e.g., personnel), financial markets (e.g., a stock market or commodity market), 10 individual company stocks and/or bonds, etc.

Brief Description of the Drawings

FIG. 1 is a block diagram of a system comprising bridge programs, and associated methodology, for generating a report by the Executive Information System (EIS) of the Systems Applications and Products (SAP) software, in accordance with embodiments of the present 15 invention.

FIG. 2 illustrates select records of a dataset, in accordance with embodiments of the present invention.

FIG. 3 is a dataset table of raw data, in accordance with embodiments of the present invention.

20 FIG. 4 is the dataset table of FIG. 3 as sorted, in accordance with embodiments of the

present invention.

FIG. 5 is the dataset table of FIG. 4 as rolled up, in accordance with embodiments of the present invention.

FIG. 6 depicts fields of an exemplary Aspect containing blocked invoices, in accordance

5 with embodiments of the present invention.

FIG. 7 depicts a detailed description of the fields of FIG. 6

FIG. 8 depicts a computer system that includes a bridge program of FIG. 1, in accordance

with embodiments of the present invention.

Detailed Description of the Invention

10 FIG. 1 is a block diagram of a system 10 comprising bridge programs, and associated

methodology, for generating a report by the Executive Information System (EIS) 12 of the

Systems Applications and Products (SAP) software, in accordance with embodiments of the

present invention. “EIS”, as referred to herein, includes not only the specific versions of EIS that

are currently operational, but also includes all future versions of EIS as well as any other

15 reporting program or module of SAP that is intended to operate on data stored in a business data

warehouse. In the preceding context, EIS may be generally viewed as an embodiment of a SAP

“business information system” having reporting capabilities in conjunction with a business data

warehouse. The present invention is relevant to any application that uses EIS, including

applications involving procurement data, such as purchase order data and invoice data. Other

20 applications may involve, *inter alia*, financial situations, human resources (e.g., personnel),

financial markets (e.g., a stock market or commodity market), individual company stocks and/or bonds, etc.

FIG. 1 shows N datasets denoted as D_1, D_2, \dots, D_N , wherein $N \geq 1$. For each i (i.e., for

$i=1, 2, \dots, N$), the dataset D_i is processed by a non-SAP bridge program P_i to generate an Aspect

5 file A_i which is sent over a data communications network W_i into an EIS 12 environment in a

form T_i . Definitionally, a data communications network comprises communication lines over

which data is transmitted from one node to another, and each said node may include, *inter alia*, a

computer, a terminal, a communication control unit, etc. A user 14 may submit a database query

to EIS 12 to obtain a report relating to data in any of T_1, T_2, \dots, T_N individually or to data in any

10 combination of T_1, T_2, \dots, T_N .

A dataset is defined as any two dimensional organization of data, or of any two-

dimensional projection of a M -dimensional organization of data wherein $M > 2$. As examples, a

dataset may be a table, database, spreadsheet, file, two-dimensional array (such as within a

computer code or database), etc. As a table, database, spreadsheet, or array, the two-

15 dimensionality of the dataset is expressed in terms of rows and columns. As a file, the two-

dimensionality of the dataset is expressed in terms of records and fields. As used herein, the

terms “row” and “record” are assumed to have the same meaning, and the terms “column” and

“field” are likewise assumed to have the same meaning. As another example, the dataset may

represent two-dimensional projection of the three-dimensional array Z_{ijk} such that the array index

20 k is constant and the array indices i and j vary. The datasets D_1, D_2, \dots, D_N may independently be

a SAP-formatted dataset or a non-SAP-formatted dataset. A SAP-formatted dataset is a dataset

that is readable by any program or module of the SAP software. A non-SAP-formatted dataset is a dataset that is not a SAP-formatted dataset. The datasets D_1, D_2, \dots, D_N have formats F_1, F_2, \dots, F_N , respectively.

The datasets D_1, D_2, \dots, D_N may be stored in any storage medium (hard disk, optical disk,

5 compact disk, magnetic tape, etc.), have any size with regard to number of bytes, be configured in accordance with any operating system on any computing platform (e.g., AIX operating system, VM operating system on VMS platform, etc.), and be located anywhere in the world or in outer

space (e.g., in the United States, Europe, South America, Asia, Africa, on a ship or submarine in the Atlantic or Pacific Ocean, in a spacecraft, on the Moon, on Mars, etc.). The datasets $D_1, D_2,$

10 \dots, D_N may have been generated through any degree of automation ranging from manual data entry to automatic generation of datasets using sophisticated software such as SAP software. As an example of manual data entry, the data for the dataset may be entered manually into a file on a floppy disk, then into a file on a workstation, then into a spreadsheet template which may be passed through the bridge program P_i to generate the Aspect file A_i ($i=1, 2, \dots, N$).

15 Each non-SAP bridge program P_i ($i=1, 2, \dots, N$) is a computer program that is not part of the SAP software and is thus not encumbered by the inefficiencies of SAP modules for processing large amounts of data. The bridge programs P_1, P_2, \dots, P_N are respectively keyed to the formats F_1, F_2, \dots, F_N of the datasets D_1, D_2, \dots, D_N for respectively generating the Aspect files A_1, A_2, \dots, A_N . Alternatively, the bridge programs P_1, P_2, \dots, P_N may be replace by a single bridge

20 program P having logical paths L_1, L_2, \dots, L_N respectively keyed to the formats F_1, F_2, \dots, F_N for respectively generating the Aspect files A_1, A_2, \dots, A_N . The bridge programs P_1, P_2, \dots, P_N (or the

single bridge program P) may be any executable form (e.g., object code, source text that can be executed by an interpreter program, a macro, etc.).

An Aspect file is a file that is readable by, and may be processed by, EIS 12 aside for a possible conversion to resolve an incompatibility between the respective operating systems and platforms in which the Aspect file and EIS 12 are configured. Thus, the Aspect file has the property that EIS 12 would be able to directly read and process the Aspect file if EIS 12 and the Aspect file were functioning on the same computing platform and with the same operating system. As an example, the Aspect files A_1, A_2, \dots, A_N may be generated by the bridge programs P_1, P_2, \dots, P_N (or the single bridge program P) using the VM operating system, while EIS 12 may be functioning in an AIX operating system environment, and file conversion of Aspect files may thus be required for operating system compatibility purposes. As another example, such Aspect file conversions may be required even if the Aspect file and EIS 12 both operate in accordance with the same operating system (e.g., AIX, UNIX, etc.). Accordingly, the Aspect files A_1, A_2, \dots, A_N , which are respectively sent over the data communications networks W_1, W_2, \dots, W_N may be respectively converted into the Temp files T_1, T_2, \dots, T_N if such conversion is necessary to resolve the operating system incompatibilities discussed *supra*. Thus, it is to be understood herein that if no such conversions are necessary then T_i and A_i ($i = 1, 2, \dots, N$) are the same Aspect file, and if such conversions are necessary then T_i is a converted form of A_i due to operating system and/or platform incompatibilities only. The term “Temp file”, as used herein, covers both of the preceding possibilities and is thus viewed as an Aspect file generated by the non-SAP bridge program of the present invention. Accordingly, the Aspect files are viewed herein as logically

readable by EIS 12 since the logical structure of the Aspect files (in terms of records and fields, or rows and columns) is adapted to be read by EIS 12.

Without the bridge programs of the present invention, a SAP dataset would normally be processed within SAP by going through SAP Logistic Information System (LIS) modules, then 5 through SAP Open Information Warehouse (OIW) modules, and finally into EIS. The preceding path of LIS → OIW → EIS is extremely time intensive and thus highly inefficient. In contrast, the bridge programs P_1, P_2, \dots, P_N (or the single bridge program P) of the present invention provide a bridge, or direct shunt from a SAP dataset or a non-SAP dataset into EIS 12 without using LIS or OIW. As will be discussed *infra*, the bridge program implements filtration and rollup 10 functionality that may drastically reduce the amount of data that is transferred from the datasets D_1, D_2, \dots, D_N into the Aspect files A_1, A_2, \dots, A_N , respectively. Thus, filtration and rollup by the bridge program significantly reduces the overall processing time for generating a report from EIS.

With filtration, the bridge program uses “selection rules” to identify “select records” of 15 the dataset. A “select record” is a record of the dataset having data to be inserted into the associated Aspect file. When filtration is used, a dataset record that is not a select record does not contribute data to the Aspect file. FIG. 2 illustrates a dataset having 15 records with records 20 1-2, 5-7, 9, and 13-14 identified by an adjacent asterisk (*) as being select records. The selection rules determine which records of the dataset are select records. A selection rule typically performs logical and/or arithmetic operations on the data in one or more fields of a records and the outcome or result of said logical and/or arithmetic operations determines whether or not the

record is a select record. An example of a selection rule is that of “date screening.” With date screening, the records of the dataset have a field that includes an effective date (e.g., the date at which a purchase order was placed, wherein the records of the dataset contain information concerning purchase orders). The “date screening” selection rule states that for a given first date

5 DATE₁ and second date DATE₂, such that DATE₁ is earlier than DATE₂, a record is identified as a select record if its effective date is not earlier than DATE₁ and not later than DATE₂.

Conversely, the record is not a select record if its effective date is earlier than DATE₁ or later

than DATE₂. Another example of a selection rule is with records that include invoice data,

6 wherein the selection rule examines the contents of a field to determine whether the invoice

10 associated with the record is a canceled invoice. A selection rule can be very complex and

involve multiple fields and/or multiple operations on the fields. Virtually any desired

combination of logical and/or arithmetic operations can be used to form a selection rule. Another

selection rule could involve testing macroscopic characteristics of a record instead of data in

fields of the record (e.g., rejecting records whose length exceed a predetermined number of

15 bytes). The selection rule may be in any programmable form such as executing a logical

statement, calling a subroutine or subprogram that performs a test that implements the selection

rule, testing data in record fields against a readout of a hardware device as a system clock or an

inline measurement instrument (e.g., a voltmeter, a pressure transducer, etc.).

FIGS. 3-5 illustrate how the bridge program performs rollup functionality. FIG. 3

20 illustrates a dataset table of raw data having 5 columns (or fields) and 22 rows (or records) as indicated. The dataset table of FIG. 3 stores purchase order records of a paint distributor

company having three divisions (100, 200, 300) and the company purchases paint for distribution to customers in two paint colors (white and blue). The fields are Purchase Order No., Vendor, Division, Color, and Buyer. Vendor (A, B, or C) denotes from whom the paint was to be purchased, and Buyer denotes the person in the company who made the purchase order. The 5 dataset table of FIG. 3 constitutes the select records of a dataset of raw records if filtration has been performed, or the entire dataset of raw records if filtration has not been performed.

The Division and Color fields collectively constitute a “keygroup” which is used for

sorting and rolling up. A keygroup generally comprises one or more fields, and said fields of the 10 keygroup are not required to be contiguously distributed along the record length. Each combination of data of the keygroup of each record of FIG. 3 is called a “keyvalue.” Each unique combination of data in the keygroups of the records of FIG. 3 is called a “rollup 15 keyvalue.” FIG. 3 shows 6 rollup keyvalues, namely: (100, blue); (100, white); (200, blue); (200, white); (300, blue); (300, white).

The 6 rollup keyvalues are clearly seen after a sort is performed on the table of FIG. 3,

15 using the first keygroup component (Division) as the primary sort key, and the second keygroup component (Color) as the secondary sort key. FIG. 4 shows the result of the aforementioned sort with an organized arrangement of the 6 rollup keyvalues.

FIG. 5 shows the table of FIG. 4 as “rolled up.” The operation of rolling up comprises

eliminating records of FIG. 4 that have redundant keyvalues, so that the records of the rolled up 20 table of FIG. 5 each have a rollup keyvalue that uniquely identifies each record and thus distinguishes each record from every other record. The rollup table of FIG. 5 has several

features. First, each record of FIG. 5 has a unique rollup keyvalue, as discussed *supra*. Second, the Purchase Order No. field and the Buyer field were eliminated, because these eliminated fields are not needed for subsequent data processing by the EIS 12 of FIG. 1. Third, the Vendor field is retained in FIG. 5, wherein the Vendor value appearing in FIG. 5 is the Vendor value in those records of FIG 4 that are also present in FIG. 5. Even if the Vendor value is not needed for subsequent data processing by the EIS 12 of FIG. 1, it is nonetheless permissible to retain the Vendor field as is done in the rolled up table of FIG. 5. Accordingly, an absence in FIG. 5 of some of the Vendor values that appears in FIGS. 3 and 4 will not adversely impact subsequent data processing by the EIS 12 of FIG. 1. Alternatively, the Vendor field could have been eliminated from the rolled up table of FIG. 5 as were the Purchase Order No. field and Buyer field. Note that for the records in FIG. 4 having redundant keyvalues, the first-appearing such record for each rollup keyvalue has been retained for placement in FIG. 5. Fourth, an added Quantity field in FIG. 5 stores the number of records in FIG. 4 having the same rollup keyvalue. The records of FIG. 5 are called “rollup records,” and the table of FIG. 5 exemplifies the Aspect file having the rollup records. Note that the rollup records in the generated Aspect file of FIG. 5 are sorted with respect to the keygroup (i.e., with respect to keyvalues of the keygroup).

In summary, given a dataset of raw records, generating the Aspect file having rollup records comprises rolling up a “portion” of the dataset of raw records with respect to the keygroup (e.g., the keygroup Division, Color in the preceding example), which 1) eliminates raw records having redundant keyvalues; and 2) includes within each rollup record of the Aspect file a “rollup” field and “quantity” field, wherein the rollup field stores a rollup keyvalue of the

keygroup, and wherein the quantity field stores the number of dataset records in FIG. 4 that have the same rollup keyvalue. Note that said “portion” of the dataset of raw records constitutes the select records of the dataset of raw records if filtration has been performed, or the entire dataset of raw records if filtration has not been performed.

5 The preceding discussion described filtration and rolling up for using a single dataset to

generate a single Aspect file. The following discussion describes analogous notation for

applying filtration and rolling up to the N datasets D_1, D_2, \dots, D_N ($N \geq 2$) in the generation of the N

Aspect files A_1, A_2, \dots, A_N , respectively, such that the Aspect files A_1, A_2, \dots, A_N have rollup

records $[R]_1, [R]_2, \dots, [R]_N$. The datasets D_1, D_2, \dots, D_N having a common keygroup. For $i = 1, 2,$

10 ..., and N, the non-SAP bridge program is executed to identify select records $[S]_i$ of the dataset

D_i , in accordance with selection rules applied to D_i . The selection rules may be the same for each

of the N datasets, or may differ for each of the N datasets. The rollup records $[R]_i$ corresponding

to $[S]_i$ have a rollup field and a quantity field. The rollup field stores a rollup keyvalue of the

select records $[S]_i$, and the quantity field stores the number of select records of $[S]_i$ that have the

15 same rollup keyvalue.

Returning to FIG. 1, it is recalled that the Aspect files A_1, A_2, \dots, A_N are transmitted to

EIS 12 where the Aspect files are the Temp files T_1, T_2, \dots, T_N , respectively. The Aspect files

and the Temp files are the same structurally (i.e., the same as to records and fields and data

therein), as explained *supra*. Thus, the Aspect files and the Temp files have the same rollup

20 records. The user 14 may make a query, such as to sum over the quantity field for a subset of the

rollup records of one or more of the Temp files. The query may be any database query that is

intended to access the Temp files. The pertinent subset of the rollup records is determined by the query itself, since the query limits the scope of the data that is required to return a result of the query, as is typical with database queries. The query is executed by a SAP module, namely EIS, in the SAP computing environment. Such execution by EIS returns a result of the query to the user 14.

5 Filtration and rollup, as implemented by the bridge programs P_1, P_2, \dots, P_N (or the bridge program P) may drastically reduce the amount of data that is transferred from the datasets D_1, D_2, \dots, D_N into the Aspect files A_1, A_2, \dots, A_N , respectively (see FIG. 1). Thus, only a small percentage of data from the raw datasets may propagate to the Temp files at EIS 12, resulting in 10 substantial reductions in processing time.

10 The bridge program is further adapted to generate a “trace file” for debugging purposes. For example, if the user 14 in FIG. 1 makes a query to EIS 12 and notices an anomalous result, such as an invoice that is outstanding for 9999 days (i.e., 27 years), then the user 14 would suspect that an error has occurred and would desire to find the source of the error. The rolling 15 up process, however, may result in eliminating fields (e.g., invoice number, invoice date, etc.) of the raw dataset in the Aspect file, wherein said fields may be required or beneficial for tracing back to the raw dataset where the error may have originated. Fortunately, the trace file of the present invention makes it possible to trace back from the Aspect file to the associated raw dataset. In particular, the trace file includes a representative rollup keyvalue of the keygroup that 20 is stored in the Aspect file and is common to the raw dataset. The trace file also includes a pointer that points to a portion of the raw dataset that is correlated with the representative rollup

keyvalue. For example, the pointer may be a representative invoice number or a representative invoice date. Accordingly, the information in the trace file may not point to exactly where in the raw dataset the problematic invoice is, but will nonetheless point to a nearby portion of the raw dataset. Thus, the nearby portion of the raw dataset is effectively correlated with the 5 representative rollup keyvalue through the pointer.

FIG. 6 lists the fields of an exemplary Aspect file that contains blocked invoices, in accordance with embodiments of the present invention. A blocked invoice is an invoice that has been submitted by an entity (e.g., company) for payment by a supplier, but the invoice contains 10 information (relating to price, quantity, taxes, receipt of goods, etc.) that is inconsistent with an associated purchase order previously issued by the entity to the supplier. The blocked invoice may be flagged for manual attention to determine the reason for the inconsistency. In FIG. 6, the “Data Type” includes a specification of the number of characters in the field (e.g., 003, 001, etc.), and the “Source” is a table name or file name identifying where the data for the field is located. 15 The detailed description of the fields listed in FIG. 6 is presented in FIG. 7.

FIG. 8 illustrates a computer system 90 for storing and executing the bridge program P_1 , P_2 , ..., or P_N of FIG. 1, in accordance with embodiments of the present invention. The computer system 90 comprises a processor 91, an input device 92 coupled to the processor 91, an output device 93 coupled to the processor 91, an output port 98 coupled to the processor 91, and 20 memory devices 94 and 95 each coupled to the processor 91. The input device 92 may be, *inter alia*, a keyboard, a mouse, etc. The output device 93 may be, *inter alia*, a printer, a plotter, a computer screen, a magnetic tape, a removable hard disk, a floppy disk, etc. The memory

devices 94 and 95 may be, *inter alia*, a hard disk, a dynamic random access memory (DRAM), a read-only memory (ROM), etc. The memory device 95 includes a bridge program 97 for generating Aspect files. The processor 91 executes the bridge program 97. The memory device 94 includes input data 96. The input data 96 includes input (e.g., datasets, control cards or files, etc.) required by the bridge program 97. The output device 93 displays printed output from the bridge program 97. The output port 98 interfaces a data communications network (see W_1 , W_2 , ..., or W_N of FIG. 1) for sending the Aspect files to the EIS 12 environment (see FIG. 1).

5 While FIG. 8 shows the computer system 90 as a particular configuration of hardware and

10 software, any configuration of hardware and software, as would be known to a person of ordinary skill in the art, may be utilized for the purposes stated *supra* in conjunction with the particular computer system 90 of FIG. 8. For example, the memory devices 94 and 95 may be portions of a single memory device rather than separate memory devices.

15 While embodiments of the present invention have been described herein for purposes of

illustration, many modifications and changes will become apparent to those skilled in the art. Accordingly, the appended claims are intended to encompass all such modifications and changes as fall within the true spirit and scope of this invention.